

A general moment expansion method for stochastic kinetic models

Angelique Ale,^{1, a)} Paul Kirk,^{1, a)} and Michael P.H. Stumpf^{1, b)}

Division of Molecular Biosciences, Theoretical Systems Biology Group, Imperial College London, London SW7 2AZ, United Kingdom

(Dated: 26 March 2013)

Moment approximation methods are gaining increasing attention for their use in the approximation of the stochastic kinetics of chemical reaction systems. In this paper we derive a general moment expansion method for any type of propensities and which allows expansion up to any number of moments. For some chemical reaction systems, more than two moments are necessary to describe the dynamic properties of the system, which the linear noise approximation (LNA) is unable to provide. Moreover, also for systems for which the mean does not have a strong dependence on higher order moments, moment approximation methods give information about higher order moments of the underlying probability distribution. We demonstrate the method using a dimerisation reaction, Michaelis-Menten kinetics and a model of an oscillating p53 system. We show that for the dimerisation reaction and Michaelis-Menten enzyme kinetics system higher order moments have limited influence on the estimation of the mean, while for the p53 system, the solution for the mean can require several moments to converge to the average obtained from many stochastic simulations. We also find that agreement between lower order moments does not guarantee that higher moments will agree. Compared to stochastic simulations our approach is numerically highly efficient at capturing the behaviour of stochastic systems in terms of the average and higher moments, and we provide expressions for the computational cost for different system sizes and orders of approximation. We show how the moment expansion method can be employed to efficiently quantify parameter sensitivity. Finally we investigate the effects of using too few moments on parameter estimation, and provide guidance on how to estimate if the distribution can be accurately approximated using only a few moments.

PACS numbers: 82.39.-k, 87.10.Mn, 82.20.Fd, 87.16.A-

I. INTRODUCTION

Cellular behaviour is shaped by molecular process that can be described by systems of chemical reactions between different molecular species. At the macroscopic scale, the dynamics of these processes are frequently described in terms of mean concentrations of species using deterministic mass action kinetics (MAK). The deterministic solution, however, does not always capture the essential kinetics of the chemical system accurately, because it excludes stochastic effects.

The stochastic kinetics of chemical reaction systems are captured by the chemical master equation (CME), a probabilistic description of the system (also known as Kolmogorov's forward equation)^{1,2}. The CME is a set of differential or difference equations that capture how the probability over the states (e.g. the abundances of the different molecular species) evolves over time. The CME offers an exact description of a system's dynamics but can only be solved analytically for very simple systems. Exact single realisations of the CME, corresponding e.g. to observing processes inside a single cell, can be obtained numerically using for example Gillespie's Stochastic Simulation Algorithm (SSA)³. The SSA is a discrete simulation method that proceeds by randomly

selecting a reaction that occurs at each subsequent time step, according to the probability of that reaction occurring next. This method is associated with considerable computational cost, that increases dramatically with the size of the model, which can make it infeasible to comprehensively characterize large systems through simulations.

A number of numerical methods have been developed to approximate the solution of the CME for more complex systems, including methods that approximate the CME by describing the probability distribution in terms of its moments⁴⁻¹¹. When only the mean (the first moment) is taken into account, the moment expansion reduces to the MAK description. In the linear noise approximation (LNA), the CME is approximated by taking into account the mean (the first moment), and the variance and covariance (second central moments) of the distribution, whereby the second central moments are decoupled from the mean¹²⁻¹⁵. This is valid in the limit of large volumes and molecule numbers, or when the system consists of first-order reactions only.

For smaller systems and more complex reactions, a number of different approaches have been developed that aim to capture the temporal changes in coupled moments. These expansions can be performed in terms of the moments or the central moments about the mean; the conversion from molecule numbers can be kept implicit in the rate parameters or made more explicit from the outset. Previous developments focussed on expansion methods for polynomials^{4,16,17}, or expansions up to a limited order. Gillespie et al. developed an expansion method in terms of moments for polynomial functions⁴;

^{a)}These authors contributed equally.

^{b)}To whom correspondence should be addressed:
a.ale@imperial.ac.uk, m.stumpf@imperial.ac.uk

this was later generalized to include rational functions as well¹⁸. Other alternative methods make the dependence on system size more obvious and include so-called Mass Fluctuation Kinetics (MFK)⁵, and the 2MA and 3MA approaches¹⁹. In the Mass Fluctuation Kinetics approach the expansion is done in terms of central moments up to third order for first and second order rate equations, and up to second order for general rate equations. In a recent paper²⁰ it was shown that expansion up to three moments tends to deliver more accurate results than expansion up to only second order; in particular the variances are improved by going to higher moments. While the MFK, 2MA and 3MA approaches include expressions for the first, second and third central moments, no general formulation exists in these frameworks to generate higher order central moments in an automatic way.

In this paper we derive a general method for expanding the CME in terms of its central moments about the mean, which does not make extraneous assumptions about the form of the reactions, and that can be evaluated up to any number of moments; in our exposition we follow the notation used in Gillespie⁴. The new method described here non-trivially generalizes the work of Gomez et al.⁵: in particular we are able to generate arbitrary order higher central moments automatically and in a computationally efficient manner; combining computer algebra systems with numerical simulation engines does allow us to tackle problems that stymy e.g. the linear noise approximation or conventional (3rd order) MFK approaches. The interplay between noise and non-linear dynamics can give rise to very complicated behaviour, and only a handful of systems have been considered. The moments of the distribution described by the CME reflect this intricate relationship between stochasticity and non-linearity and we spend some time discussing this for a typical non-linear biomolecular system. Thus while going to arbitrary moments is straightforward in our framework, we focus our discussion on the lower (up to sixth order) moments.

Our expansion is based on two successive Taylor expansions that allow us to express the CME of the moment generating function in computationally convenient form; and we truncate the system by setting higher order terms in the Taylor expansions to zero. The outlined method could fit into a general framework for parameter inference based on maximum entropy distributions derived from the calculated moments. We furthermore demonstrate that parameter sensitivity analyses may be performed naturally and efficiently in this framework: we can consider the rate of change of the moments with respect to the parameters, which also allows us to study the factors underlying cell-to-cell variability. We illustrate the general method using three molecular reaction systems: a simple dimerisation reaction, which allows for a detailed investigation; Michaelis-Menten enzyme kinetics; and the p53 system, an oscillating tumour suppressor system²¹.

II. MOMENT EXPANSION METHOD

We consider a system with N different molecular species (X_1, \dots, X_N) that are involved in L chemical reactions with reaction rates k_l ,

$$\underline{s}_1 X_1 + \dots + \underline{s}_N X_N \xrightarrow{k_l} \bar{s}_1 X_1 + \dots + \bar{s}_N X_N, \quad (1)$$

with \underline{s}_i and \bar{s}_i the number of molecules of type X_i before and after the reaction, respectively. The time evolution of the system's state is described by the chemical master equation (CME),

$$\frac{dP(\mathbf{x})}{dt} = \sum_l P(\mathbf{x} - \mathbf{s}_l) a_l(\mathbf{x} - \mathbf{s}_l) - P(\mathbf{x}) a_l(\mathbf{x}), \quad (2)$$

in which $a_l(\mathbf{x})$ are the propensity functions, with $a_l(\mathbf{x})dt$ defined as the probability of reaction l occurring in an infinitesimal time interval dt when the number of molecules in the system is \mathbf{x} , $P(\mathbf{x})$ the probability that the system contains \mathbf{x} molecules and $\mathbf{s}_l = \bar{\mathbf{s}}_l - \underline{\mathbf{s}}_l$.

We start the derivation of the moment expansion method by deriving a moment generating function from the CME. In general, the moment generating function of a random variable \mathbf{x} is defined as¹

$$m(\theta, \mathbf{x}) = \sum_{\mathbf{x}} e^{\theta \mathbf{x}} P(\mathbf{x}). \quad (3)$$

The moments, $\langle \mathbf{x}^{\mathbf{n}} \rangle$, with $\mathbf{x}^{\mathbf{n}} = x_1^{n_1} \dots x_d^{n_d}$, of the probability distribution can be found by taking the n -th order derivatives of the moment generating function with respect to θ . The first moment is, of course, equal to the mean, $\mu = \langle x \rangle$. The variance, skewness (related to asymmetry of the distribution) and kurtosis (related to the chance of outliers) can be derived from the central moments about the mean, $\langle (\mathbf{x} - \mu)^{\mathbf{n}} \rangle$.

Using the CME we can write the time dependent moment generating function²²

$$\frac{dm}{dt} = \sum_l \left[(e^{\theta \mathbf{s}_l} - 1) \sum_{\mathbf{x}} e^{\theta \mathbf{x}} P(\mathbf{x}) a_l(\mathbf{x}) \right] \quad (4)$$

The time evolution of the mean concentration μ_i of species X_i can be obtained by taking the first derivative of Eq. 4 with respect to θ_i and subsequently setting θ to zero,

$$\frac{d\mu_i}{dt} = \frac{d}{d\theta_i} \frac{dm}{dt} \Big|_{\theta=0} = \sum_l s_{li} \langle a_l(\mathbf{x}) \rangle \quad (5)$$

This expression can be evaluated by expanding $a_l(\mathbf{x})$ in a Taylor expansion about the mean,

$$\frac{d\mu_i}{dt} = S \left[\sum_l \sum_{n_1=0}^{\infty} \dots \sum_{n_d=0}^{\infty} \frac{1}{\mathbf{n}!} \frac{\partial^{\mathbf{n}} a_l(\mathbf{x})}{\partial \mathbf{x}^{\mathbf{n}}} \Big|_{\mathbf{x}=\mu} \mathbf{M}_{\mathbf{x}^{\mathbf{n}}} \right], \quad (6)$$

where S is the stoichiometry matrix,

$$\partial^{\mathbf{n}} = \partial^{n_1 + \dots + n_d}$$

$$\mathbf{n}! = n_1! \dots n_d!$$

$$\mathbf{M}_{\mathbf{x}^{\mathbf{n}}} = M_{x_1^{n_1}, \dots, x_d^{n_d}}$$

$$\partial \mathbf{x}^{\mathbf{n}} = \partial x_1^{n_1} \dots \partial x_d^{n_d},$$

and we have substituted the central moments around the mean for the expected values of the expansion terms

$$M_{x_1^{n_1}, \dots, x_d^{n_d}} = \langle (x_1 - \mu_1)^{n_1} \dots (x_d - \mu_d)^{n_d} \rangle.$$

The first central moment $\langle \mathbf{x} - \boldsymbol{\mu} \rangle$ is zero. Higher order central moments can be derived from the moments, for example the covariance between x_1 and x_2 can be written as

$$\sigma_{x_1 x_2}^2 = \langle (x_1 - \mu_1)(x_2 - \mu_2) \rangle = \langle x_1 x_2 \rangle + \mu_1 \mu_2 - \langle x_1 \rangle \langle \mu_2 \rangle - \langle x_2 \rangle \langle \mu_1 \rangle. \quad (7)$$

In general the relation between the central moments, $\mathbf{M}_{\mathbf{x}^n}$, and the moments, μ^n , can be formulated as

$$\mathbf{M}_{\mathbf{x}^n} = \sum_{k_1=0}^{n_1} \dots \sum_{k_d=0}^{n_d} \binom{\mathbf{n}}{\mathbf{k}} (-1)^{(\mathbf{n}-\mathbf{k})} \underbrace{\mu^{(\mathbf{n}-\mathbf{k})}}_{\alpha} \underbrace{\langle \mathbf{x}^{\mathbf{k}} \rangle}_{\beta}, \quad (8)$$

where

$$\begin{aligned} (-1)^{(\mathbf{n}-\mathbf{k})} &= (-1)^{(n_1-k_1)} \dots (-1)^{(n_d-k_d)} \\ \mu^{(\mathbf{n}-\mathbf{k})} &= \mu_1^{(n_1-k_1)} \dots \mu_d^{(n_d-k_d)} \\ \mathbf{x}^{\mathbf{k}} &= x_1^{k_1} \dots x_d^{k_d} \\ \binom{\mathbf{n}}{\mathbf{k}} &= \binom{n_1}{k_1} \dots \binom{n_d}{k_d}. \end{aligned}$$

We obtain the time evolution equations of the central moments by taking the time derivative of Eq. 8,

$$\frac{d\mathbf{M}_{\mathbf{x}^n}}{dt} = \sum_{k_1=0}^{n_1} \dots \sum_{k_d=0}^{n_d} \binom{\mathbf{n}}{\mathbf{k}} (-1)^{(\mathbf{n}-\mathbf{k})} \left[\alpha \frac{d\beta}{dt} + \beta \frac{d\alpha}{dt} \right] \quad (9)$$

The term α makes the time evolution of the central moments a function of Eq. 6,

$$\frac{d\alpha}{dt} = \sum_{i=1}^N (n_i - k_i) \mu_i^{-1} \alpha \frac{d\mu_i}{dt} \quad (10)$$

The term β gives rise to mixed moments, and the derivative of β with respect to time yields the time evolution equations for the mixed moments. Therefore we also need to include the time derivatives of the mixed moments in our system of equations. We obtain the time derivative of the term $\langle \mathbf{x}^{\mathbf{k}} \rangle$ by taking higher order derivatives of the moment generating function Eq. 4, resulting in

$$\frac{d\beta}{dt} = \sum_{e_1=0}^{k_1} \dots \sum_{e_d=0}^{k_d} \mathbf{s}^e \binom{\mathbf{k}}{\mathbf{e}} \underbrace{\langle \mathbf{x}^{(\mathbf{k}-\mathbf{e})} a(x) \rangle}_{\langle F \rangle}, \quad (11)$$

where

$$\begin{aligned} \mathbf{s}^e &= s_1^{e_1} \dots s_d^{e_d} \\ \mathbf{x}^{(\mathbf{k}-\mathbf{e})} &= x_1^{k_1-e_1} \dots x_d^{k_d-e_d}. \end{aligned}$$

By expanding the individual terms of the resulting expressions in a second Taylor expansion, the time evolution of the mixed moments becomes a function of the

central moments; the time evolution of the central moments remains a function of the central moments alone,

$$\langle F \rangle = \sum_{n_1=0}^{\infty} \dots \sum_{n_d=0}^{\infty} \frac{1}{\mathbf{n}!} \frac{\partial^{\mathbf{n}} F(x)}{\partial \mathbf{x}^{\mathbf{n}}} \bigg|_{x=\boldsymbol{\mu}} \mathbf{M}_{\mathbf{x}^{\mathbf{n}}}. \quad (12)$$

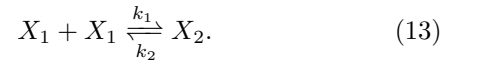
When the model under investigation is non-linear, each central moment will depend on a higher order central moment, which may itself also depend on higher order moments; hence the number of equations we would need to include is in principle infinite. To overcome this, we can obtain a closed set of equations by evaluating the time evolution equations for ν moments and truncating the Taylor series after the ν th order, thereby setting all higher order central moments equal to zero. By truncating the Taylor expansion (i.e. setting terms of the Taylor expansion corresponding to $\sum n_i > \nu$ to 0), the equations are only dependent on the central moments up to the selected order ν . Alternatively, the set of equations could be closed using moment closure techniques based on common expressions for well known probability distributions¹⁸. In this paper we will use truncation as well as closure based on a Gaussian distribution.

III. RESULTS

We illustrate the approach in a range of applications that serve to highlight both the basic properties of the moment expansion method as well as the advantages this method offers in situations where other methods^{15,23,24} tend to fail.

A. Dimerisation

We first illustrate the moment approximation method using a simple dimerisation reaction²⁵. The system describes the reversible formation and disintegration of a dimeric molecule,



The system can be written in terms of two propensities

$$a_1 : k_1 x_1(x_1 - 1); \quad a_2 : k_2 x_2, \quad (14)$$

and the stoichiometry matrix

$$S = \begin{bmatrix} -2 & 2 \\ 1 & -1 \end{bmatrix}, \quad (15)$$

where the columns correspond to reactions and the rows to variables.

The exact kinetics of the system can be straightforwardly simulated using the Stochastic Simulation Algorithm (SSA)³. One realisation of the SSA is equivalent to observing the kinetics of the system inside a single cell (Figure 1a), whereas the average over many realisations mimics the observation of the average kinetics for a large number of cells (Figure 1b, $n=100,000$). The system

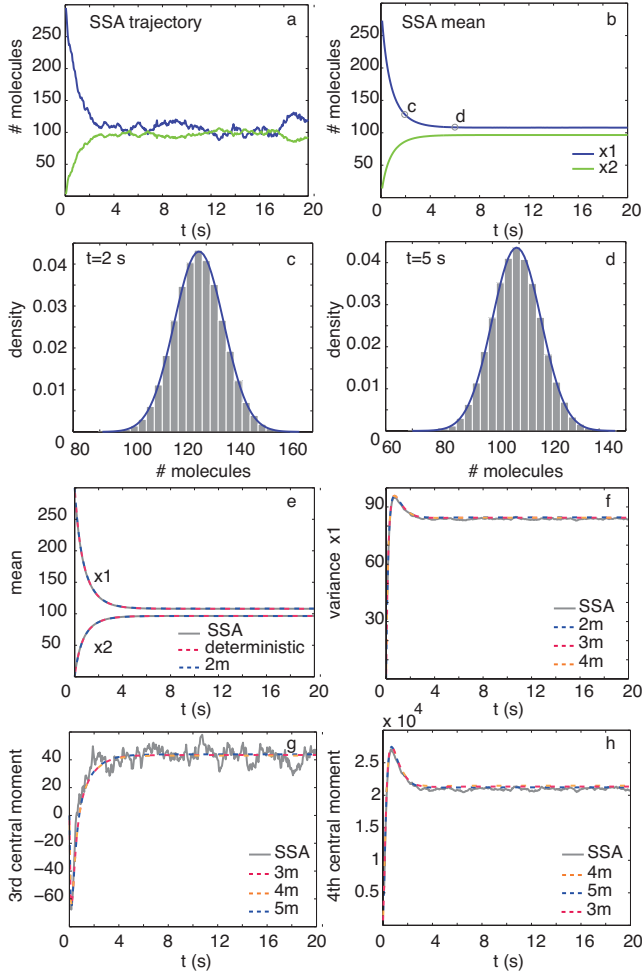


FIG. 1. Study of dimerization system, initial values $\mathbf{x} = [301, 0]$, parameters $\mathbf{k} = [1.66 \cdot 10^{-3}, 0.2]$. a) Single SSA realisation. b) Average of 100,000 SSA simulations. c-d) Histogram of SSA runs (grey bars) and probability density of normal distribution (blue line) calculated from mean and variance of SSA runs corresponding to points c and d in figure (b). e) Mean for both variables, calculated using SSA, the moment approximation using 1 moment (deterministic) and two central moments (2m). f) Variance of x_1 calculated using SSA, two central moments (2m), three central moments with Gaussian closure (3m) and four central moments (4m). g) Third central moment calculated using SSA and moment approximation method, fourth central moment calculated using SSA and moment approximation method.

reaches a stationary state after about 4 seconds. In Fig. 1c-d we plotted histograms of the SSA at two time points that correspond to different regimes in the trajectory: transient state with decreasing molecule number (Fig. 1c) and stationary state (Fig. 1d). We calculated the normal probability density function from the mean and variance of the SSA distribution at those two time points (blue line), ignoring higher order moments. The normal probability density functions fit the histograms very well, which indicates that the distribution of molecule numbers is approximately normal over the time course, and

TABLE I. Error between mean, second and third central moment calculated with SSA and approximation methods for the dimerization system.

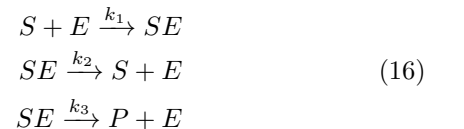
ϵ [%]	deterministic	2m	3m	4m	5m	6m
μ_1	0.300	0.0545	0.0546	0.0546	0.0546	0.0546
$M_{x_1}^2$		0.961	0.803	0.804	0.801	0.803
$M_{x_1}^3$			18.3	18.1	18.2	18.2
$M_{x_1}^4$				2.39	1.33	1.83

higher order moments would have limited influence on describing the kinetics of the mean for this system.

Figure 1d shows the mean molecule numbers calculated with the moment approximation method compared to the SSA results. The results for the deterministic (including only the mean) as well as the two moment approximation are approximately equal to the means calculated with the SSA. We compare the higher order central moments calculated with the general moment approximation method described above with the results from the SSA simulations in Figures 1f-h. In the 3-moment approximation we closed the equations using the Gaussian probability distribution (setting the fourth cumulant equal to zero²⁰), while in the other approximations we truncated higher order moments. The approximated moments are close to the exact moments calculated from the SSA, which is also clear from the errors displayed in Table I, calculated as $\epsilon = (1/N) \sum_n \sqrt{((M_{SSA} - M_{ma})/M_{SSA})^2}$ with M_{SSA} the moment or central moment calculated based on the SSA, M_{ma} the corresponding value calculated with the moment approximation, and N the number of time points taken into account. The larger error for the mean when using the deterministic approach is due to small differences in the first part of the trajectory, the decreasing part, where the contributions of the higher order central moments are largest. The larger errors calculated for the third central moment are due to the fluctuations that are still present in the third central moment trajectory calculated based on 100,000 SSA runs. Increasing the number of SSA runs would reduce this effect.

B. Michaelis-Menten enzyme kinetics

We next look at Michaelis-Menten enzyme kinetics, where an enzyme, E , and substrate, S , first bind to form a complex, SE ; following this, the complex can dissociate into the original components S and E , or S can be converted into the product, P ,



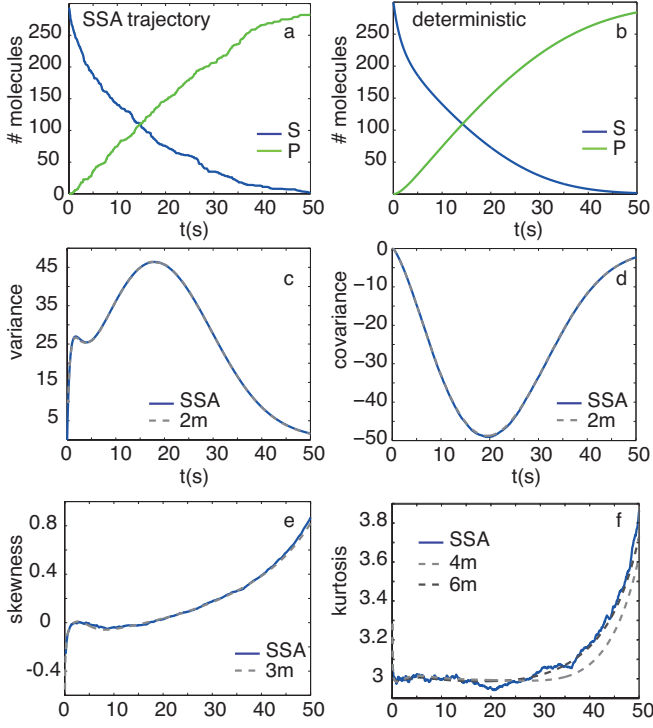


FIG. 2. Study of Michaelis-Menten kinetics, with parameters $\mathbf{k} = [1.66 \cdot 10^{-3}, 1 \cdot 10^{-4}, 0.01]$, and initial conditions, $S(0) = 301$, $P(0) = 0$, and $E(0) = 120$. (a) Single SSA realisation. Trajectories calculated using (b) moment approximation including only the mean (deterministic). (c-d) Variance of S and covariance between S and P calculated using SSA and approximation using 2 moments. (e) Skewness of S calculated using SSA and approximation with 3 moments, (f) Kurtosis calculated using SSA and approximation up to 4 and 6 moments.

The system is often reduced to a system of two variables (S and P)²⁵, with three reaction propensities

$$\begin{aligned} a_1 &: k_1 S[E(0) - S(0) + S + P]; \\ a_2 &: k_2 [S(0) - (S + P)]; \\ a_3 &: k_3 [S(0) - (S + P)] \end{aligned}$$

and the stoichiometry matrix

$$S = \begin{bmatrix} -1 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}. \quad (17)$$

One trajectory calculated with Gillespie's Stochastic Simulation Algorithm is shown in Figure 2a, and the mean calculated by solving the ODE system using the deterministic representation of the system is displayed in Figure 2b. For this system the deterministic representation is very close to the stochastic solution. The variance of the substrate and the covariance between the substrate and the product (Figure 2c-d) calculated based on 100,000 SSA runs can be closely approximated using the moment approximation method expanded up to two

moments. Figure 2e shows the skewness of the distribution over time, calculated as

$$\gamma = \frac{\langle (x_1 - \mu_1)^3 \rangle}{\sigma_{11}^3} = \frac{\langle (x_1 - \mu_1)^3 \rangle}{\langle (x_1 - \mu_1)^2 \rangle^{3/2}}. \quad (18)$$

For a normal distribution the skewness is zero. The skewness is approximated well using the moment approximation method up to three moments. The kurtosis, given by

$$\gamma_2 = \frac{\langle (x_1 - \mu_1)^4 \rangle}{\sigma_{11}^4} = \frac{\langle (x_1 - \mu_1)^4 \rangle}{\langle (x_1 - \mu_1)^2 \rangle^{4/2}}, \quad (19)$$

indicates the thickness of the tails of the probability distribution, relating to the probability of outliers. For a normal distribution, the kurtosis is equal to 3. When four moments are used to approximate the system, the approximation of the kurtosis that we obtain from the SSA is not as close as when also higher moments, here six moments, are included in the calculation. This illustrates that agreement between lower-order moments does not guarantee that higher-order moments will also agree. This problem is likely exacerbated for more complex models, e.g. those exhibiting non-linear dynamics.

C. P53 system

Finally, we investigate the oscillatory p53 system²⁶, which consists of three proteins connected through a non-linear feedback loop: p53, precursor of Mdm2 and Mdm2. The system can be written in terms of six propensities,

$$\begin{aligned} a_1 &: k_1; & a_2 &: k_2 x \\ a_3 &: k_3 \frac{xy}{x+k_7}; & a_4 &: k_4 x \\ a_5 &: k_5 y_0; & a_6 &: k_6 y, \end{aligned} \quad (20)$$

and the stoichiometry matrix

$$S = \begin{bmatrix} 1 & -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}, \quad (21)$$

where x is the concentration of p53, y_0 the concentration of precursor of Mdm2, y is the concentration of Mdm2, k_1 is the p53 production rate, k_2 is the Mdm2-independent p53 degradation rate, k_3 the saturating p53 degradation rate, k_7 is the p53 threshold for degradation by Mdm2, k_4 is the p53-dependent Mdm2 production rate, k_5 is the Mdm2 maturation rate and k_6 is the Mdm2 degradation rate.

Figure 3a shows an SSA simulation of the model for parameters $q_1 = [90, 0.002, 1.7, 1.1, 0.93, 0.96, 0.01]$ and initial values $\mathbf{x}(0) = [70, 30, 60]$, which exhibits oscillating behaviour. Because the oscillations for different realisations of the SSA (corresponding to different cells) are stochastically out of phase, the average over 100,000 stochastic simulations shows a damped oscillation, reflecting the

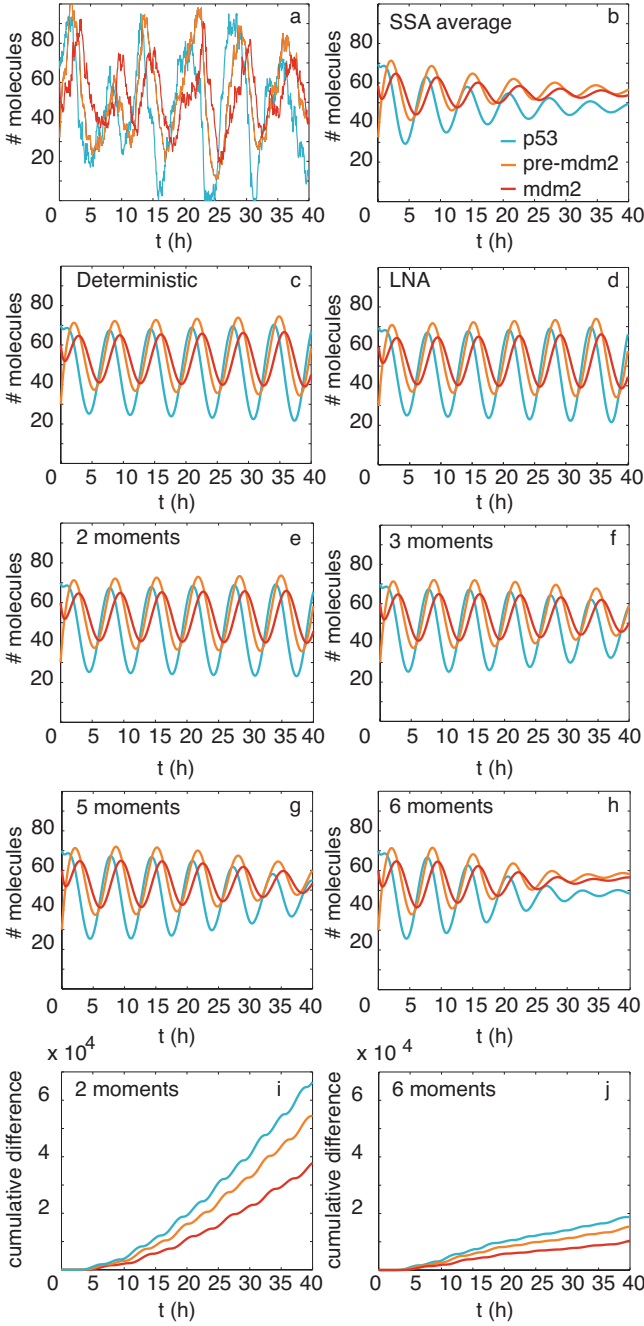


FIG. 3. Study of p53 model, parameter set $q_1 = [90, 0.002, 1.7, 1.1, 0.93, 0.96, 0.01]$. (a) Single SSA realisation. (b) Average of 100,000 SSA simulations. Trajectories calculated using (c) moment approximation including only the mean (deterministic), (d) Linear Noise Approximation, (e) mean and the variance, (f) up to three central moments, (g) five central moments and (h) six central moments. (i-j) Cumulative difference between mean trajectory calculated with SSA and trajectories calculated using (i) 2 moments and (j) 6 moments.

presence of a negative feedback loop. Figures 3c-h show the trajectories of the mean calculated with the moment approximation method. In the deterministic approximation, the oscillations are not damped but expanding,

which would indicate a positive instead of negative feedback loop. The LNA and the 2 moment approximation show the same effect. The mean calculated with the LNA is always equal to the mean calculated with the deterministic approximation because the mean is not coupled to the variance. When 3 moments are included, the system shows damped oscillations, although not as damped as the SSA trajectories. By including more moments the trajectories converge to the SSA trajectories. When 6 moments are taken into account, the trajectories calculated with the moment approximation show a similar behaviour to the trajectories calculated with the SSA. This is confirmed by the cumulative difference between the SSA trajectories and the trajectories calculated with the moment approximation shown in Figures 3i-j, which show a clear decrease in cumulative error for all variables when 6 central moments compared to 2 central moments are included in the approximation. Including more moments would improve the estimation further.

We analyze the distribution of the p53 model over the time course by looking at the central moments (Figure 4). The variance for the SSA is first increasing, then after about 20 hours it levels off. In Figure 4b we compare the variance calculated based on the SSA with the LNA and moment approximation method. When up to five central moments are included, the variance keeps increasing and does not level off. Only for the case of 6 moments does the variance reach a stable value, and even then the value is about three times higher than that predicted by the SSA simulations. Figure 4c shows the comparison of the skewness of the SSA distribution to the skewness of a normal distribution ($\gamma = 0$). For three time points where the skewness is relatively large (indicated by d, e, f) we display the histogram of the 100,000 SSA realisations together with the probability density of the normal distribution (cyan line) calculated based on the mean and variance of the SSA and the 6 moment approximation. Additionally, we plot the skew-normal distribution calculated using the mean, variance and third central moment; this is defined by the probability density function

$$f(x) = \frac{2}{\omega} \phi\left(\frac{x - \xi}{\omega}\right) \psi\left(\alpha \left(\frac{x - \xi}{\omega}\right)\right), \quad (22)$$

with ξ a location parameter and ω a scale parameter. The parameter δ can be calculated from the estimated skewness using the relation

$$|\delta| = \sqrt{\frac{\pi}{2} \frac{|\hat{\gamma}|^{2/3}}{|\hat{\gamma}|^{2/3} + ((4 - \pi)/2)^{2/3}}}. \quad (23)$$

The parameter α can be calculated from δ with $\delta = \alpha/(\sqrt{1 - \delta^2})$, and ω can be calculated from the variance using $\sigma^2 = \omega^2(1 - 2\delta^2/\pi)$. These plots confirm what we saw above, that using only the mean and variance does not capture the full distribution in this case, and also including the skewness is not enough. When comparing the skewness of the p53 distribution with the skewness of the Michaelis-Menten enzyme kinetics system (Figure

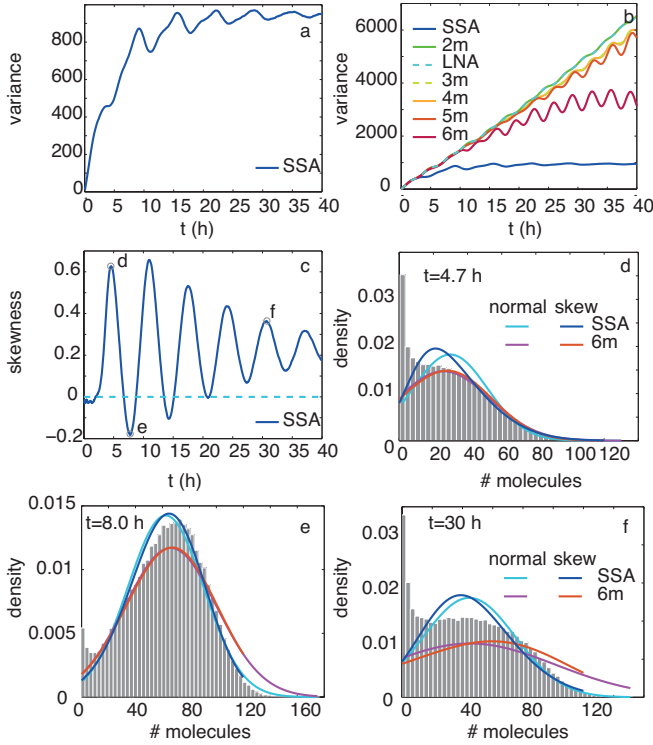


FIG. 4. Analysis of distribution for p53 model. (a) Variance calculated based on SSA runs. (b) Variance calculated with SSA, LNA and moment approximation method. (c) Skewness calculated based on SSA runs (blue line) and skewness for normal distribution (cyan dashed line). (d-f) Histograms calculated based on SSA for points d, e and f in figure (c), and probability density function of normal distribution calculated using mean and variance based on SSA (cyan line).

2e), we find that the maximum value of the skewness for both systems is approximately equal, and in both systems the skewness does not have a large effect on the mean.

We evaluate the influence of ignoring higher order moments on parameter estimation with this system by looking at 2-dimensional contour diagrams of the distance d (which is closely related to the Gaussian likelihood²⁷) calculated via

$$d = -0.5 \sum_{t=1}^T \sum_{i=1}^3 (x_i - \mu_i)^2 \quad (24)$$

where t are the time points of the trajectories and T the total number of time points included in the calculation ($T = 800$), x_i are the mean values of the SSA trajectories for the three variables in the system at timepoints t , and μ_i are the mean values of the trajectories calculated with the moment expansion method. Figures 5a-b show contour plots of the distances between the SSA trajectories and moment expansion trajectories calculated for varying values of parameters k_1 and k_3 in the p53 system (Eq. 20), Figures 5c-d show contour plots of the distances calculated for varying values of parameters k_5

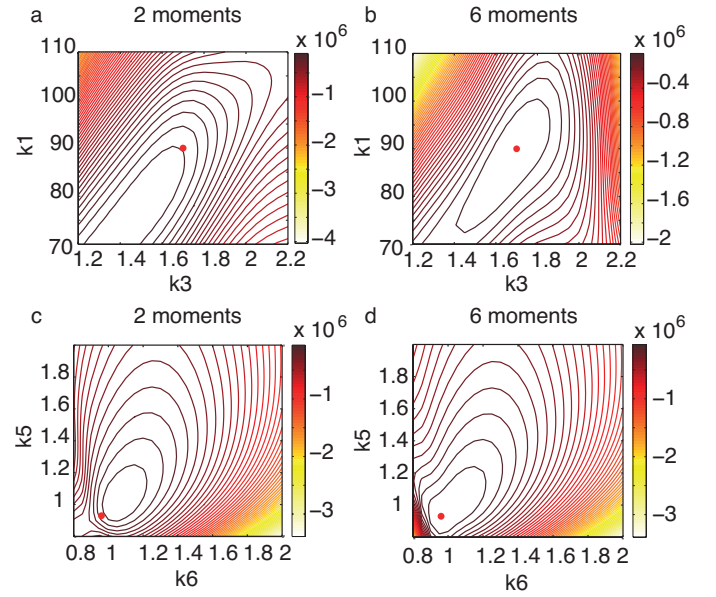


FIG. 5. Contour plots for the p53 system of the distance d between SSA trajectories and trajectories calculated with the moment expansion method. a-b) Contour for varying k_1 and k_3 , calculated using expansion up to 2 and 6 moments respectively. c-d) Contour for varying k_5 and k_6 , calculated using expansion up to 2 and 6 moments respectively. Red dots indicate parameter values used for SSA simulations.

and k_6 . The parameters used for the SSA simulations ($k_1 = 90$, $k_3 = 1.7$, $k_5 = 0.93$, $k_6 = 0.96$) are indicated with a red dot. From the contour plots corresponding to expansion up to two central moments (Figures 5a,c) we can see that the minimum distance is obtained for parameters values that are not equal to the actual parameters used for the SSA simulations. This indicates that when expansion up to two moments is used for parameter inference of the p53 system, the wrong parameters will be found. However, when six central moments are included (Figures 5b,d), the minimum distance does correspond to the parameters used for the SSA calculations.

D. Parameter Sensitivity Estimation

Assessing parameter sensitivity is a key concern when fitting any parametric model^{28,29}. Such analyses allow us to quantify how rapidly the outputs of our model change as we vary its parameters, which can provide insights into the robustness of the model and the relative influence that each parameter has upon the model's behaviour. However, sensitivity analyses of stochastic models can be difficult and/or computationally costly^{30,31}, and often involve simulating many times in order to obtain Monte Carlo estimates of sensitivity coefficients. The development of efficient methods for stochastic sensitivity analyses has been the focus of much recent research³¹⁻³³.

In the context of our proposed approach, a natural

and straightforward way to assess parameter sensitivity is to consider the rate at which the moments vary with the parameters. This motivates the calculation of simple sensitivity coefficients^{28,34} of the form $s_{ij}(t) = \frac{\partial m_i(t, \theta)}{\partial \theta_j}$, where m_i is the (estimated) i -th moment and θ_j is the j -th parameter. Within our moment approximation framework, the s_{ij} 's may either be estimated by perturbing the model's parameters and computing a finite difference approximation, or obtained automatically by employing the CVODES solver of Serban and Hindmarsh³⁵ when solving the system of ODEs (Equations (6) and (9)). In Figure 6, we reconsider the dimerisation model of Section III A. We focus upon the sensitivity of the mean and 2nd and 3rd central moments of the two molecular species to the parameter k_1 (similar results are obtained for the sensitivity to k_2). Figures 6a, d and g show how the moments estimated from 100,000 SSA runs vary when we increase the original value of k_1 by 10 percent. Figures 6b, e and h show the same for the moments estimated using our proposed approach with 6 central moments (6m). Figures 6c, f and i show sensitivity coefficients estimated from both the SSA and 6m outputs using a finite difference approximation (in the 6m case, the sensitivity coefficients may instead be obtained automatically using the CVODES solver, which yields identical results). There is generally good agreement between the coefficients estimated using the two different approaches. However, as we consider higher moments, our ability to assess sensitivity using the SSA output rapidly diminishes, since the variability caused by the change in the parameter value is overwhelmed by the variability in the estimator due to finite sample size. This may be rectified by increasing the number of SSA simulations, but at considerable computational cost. In contrast, the sensitivity coefficients associated with higher moments may still be straightforwardly calculated using the moment expansion approach (although, of course, care must be taken to ensure that appropriately many moments have been taken into account by the approximation — see Section III E).

E. Simple Heuristics for Moment Expansions

Our results for the p53 system clearly demonstrate that failure to take a sufficient number of moments into account can lead to incorrect conclusions and biased parameter estimates. Ideally we would like to know from the outset whether a deterministic approach or a two moment approximation is sufficient to capture the general statistical behaviour of the stochastic system. But without recourse to a large number of SSA runs it is impossible to predict the statistical properties of the solutions to non-linear stochastic systems. And in such cases it is generally not feasible to perform large numbers of SSA simulations and we need a different approach. We should look at the assumption made at the beginning of our derivation, where we assumed that we can approximate the propensity functions with a Taylor expansion.

For a single variable a Taylor expansion of a function $f(x)$ about the point c has the general form

$$f(x) = f(c) + \frac{f'(c)}{1!} (x - c) + \frac{f''(c)}{2!} (x - c)^2 + \dots \quad (25)$$

By taking into account only the first term, we assume that the function value in point x is the same as for point c ; by taking into account also the second term we assume that $f(x)$ can be approximated by a straight line between points x and c , etc.. Truncating the expansion at a low order will only result in a good approximation in case x is close to point c , where we have approximated the true function. In our case c is the mean, μ , implying that an approximation using a few moments will be accurate only in case all observations are close to the mean. In case it is possible to perform a single realisation of the SSA, we can assess this quality by comparing the mean calculated with the deterministic approach with the trajectory calculated with the SSA (Figure 7).

Figure 7a displays the deterministic mean of x_1 and one SSA simulation for the dimerisation system. In this case the trajectory is close to the mean over the complete time course. A single trajectory for p53, displayed in Figure 7c, compared to the deterministic result shows that in this case the distance of the trajectory from the mean is much larger. In case a single SSA realisation of the model is not possible, but experimental data are available, the distance from the mean can still be investigated in the same way. Figures 7b,d show the mean calculated using the deterministic approach together with ‘measured’ values at three time points (obtained with the SSA), repeated three times, resulting in nine data points. Also when looking at the distance from these nine points from the deterministic mean, it is clear that for the dimerisation system $(x - \mu)$ is small, whereas for the p53 system it is relatively large, indicating that a larger number of moments is necessary to capture the full distribution.

Such simple heuristics have the advantage of being computationally affordable. While inadequate at guaranteeing good performance of an expansion using any finite number of moments, we can use them to capture any gross inadequacy of a given approximation relatively reliably. Such small-scale analyses should precede or accompany moment expansions. More generally, we can consider this problem from the point of view of statistical model checking; see e.g.³⁶. But the question as to how many stochastic simulations need to be averaged over to get a good idea of the mean (or any higher moment) is challenging to answer for all but the most trivial systems³⁷.

F. Computational Complexity

The computational complexity of the moment approximation method depends on the number of variables in the system, and the number of terms that need to be

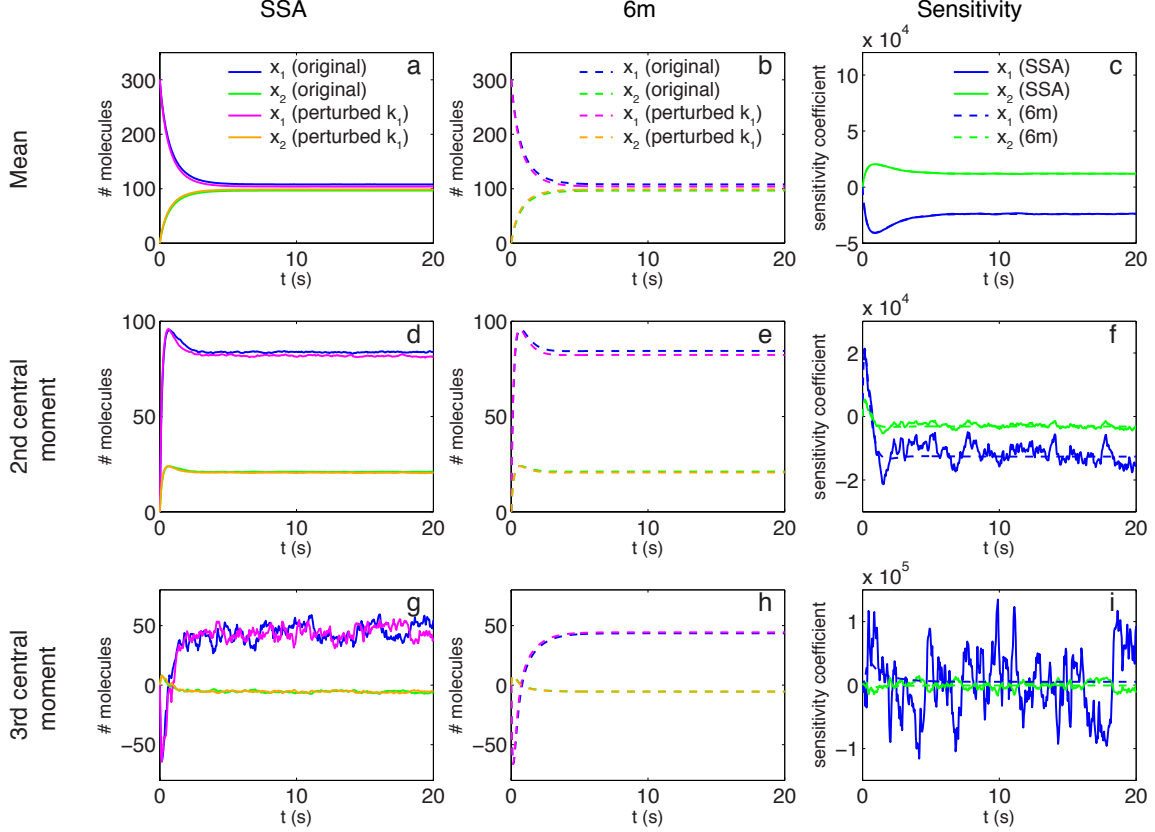


FIG. 6. Assessing the sensitivity to parameter k_1 of the dimerisation system of Section III A. Initial conditions and parameters are the same as previously, except that we additionally consider a perturbed value of $k_1 = 1.1 \times 1.66 \cdot 10^{-3}$ a) Average of 100,000 SSA simulations for the original and perturbed k_1 values. b) As (a), except the mean is estimated using the proposed moment expansion approximation with 6 moments. c) The sensitivity coefficients $\partial\mu_1(t, [k_1, k_2])/\partial k_1$ and $\partial\mu_2(t, [k_1, k_2])/\partial k_2$ (where μ_1 and μ_2 represent the means of x_1 and x_2 respectively) estimated from the SSA and 6m output via finite different approximations. d), e) and f) As a), b) and c), but for the second central moment. g), h) and i) As a), b) and c), but for the third central moment.

evaluated for each central moment. Because of the symmetry of central moments, e.g. $\langle(x_1 - \mu_1)(x_2 - \mu_2)\rangle = \langle(x_2 - \mu_2)(x_1 - \mu_1)\rangle$, there are many terms we do not need to include in the ODE representation of the system. The total number of central moment terms that could be nonzero when approximating a system with d variables and up to N moments is given by

$$N_{cm} = \frac{(N+d)!}{N!d!} - d - 1 \quad (26)$$

We subtract d terms because the first order central moments are always zero, and one term corresponding to the zeroth order central moment. For the deterministic case, the total number of ODE equations necessary to describe the system is equal to the number of variables, each term describes the mean of one of the variables. For the LNA the total number of equations is equal to the

number of equations needed for the two moment approximation. We displayed the number of central moment terms to be evaluated for systems up to 9 variables and up to 6 central moments in Table II. When two central moments are included for 2 variables, we need to evaluate two ODE equations corresponding to the mean, and three equations for the second central moments, i.e. the variance of both variables, and their covariance. From Table II we can see that the number of equations that need to be included rises quickly with the number of variables in the system.

However, in order to obtain the same information on higher order moments by simulation with the SSA, a large number of simulation runs would have to be performed; in our experience even for relatively simple systems this is of the order of $> 10^4 - 10^6$. We have already seen from the examples described above that the num-

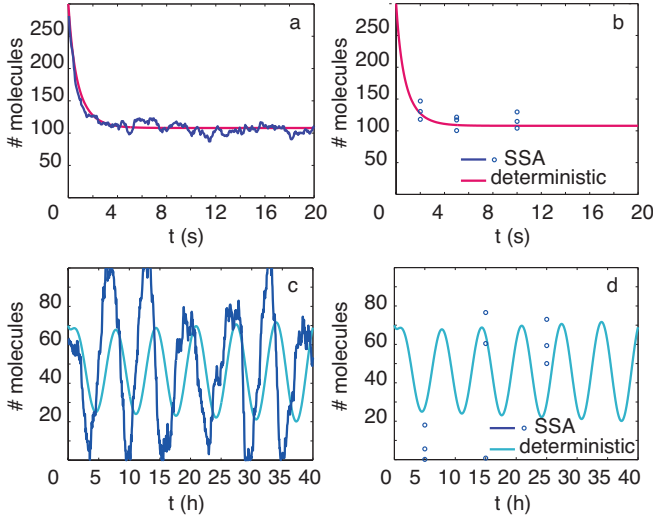


FIG. 7. Study of the deviation from the mean ($x - \mu$) (a) Deterministic mean and single SSA trajectory for dimerisation system. (b) Deterministic mean and 9 points taken from different SSA trajectories for dimerisation system. (c) Deterministic mean and single SSA trajectory for p53 system. (d) Deterministic mean and 9 points taken from different SSA trajectories for p53 system.

TABLE II. Number of potentially nonzero central moment terms to include in moment approximation for different numbers of variables (columns) and number of central moments (rows).

	N variables							
	2	3	4	5	6	7	8	9
$2m$	3	6	10	15	21	28	36	45
$3m$	7	16	30	50	77	112	156	210
$4m$	12	31	65	120	203	322	486	705
$5m$	18	52	121	246	455	784	1278	1992
$6m$	25	80	205	456	917	1708	2994	4995

ber of simulations necessary to obtain a smooth estimate of the higher order central moments increases with each order. For example, the third central moment in the dimerisation system calculated from 100,000 simulations still displays fluctuations that can only be expected to smooth out when more simulations were performed.

IV. CONCLUSION

In this paper we have described a general moment expansion method for approximating the time evolution of stochastic kinetic systems. We have shown that taking into account more moments improves the estimate of the mean and the higher order moments. In case the deterministic approach delivers an accurate estimation of the mean, expanding the higher order moments still gives ad-

ditional information on the variance, skewness and kurtosis of the distribution of the variables. Even for very simple systems, e.g. the dimerisation system, we find that higher moments obtained from averaging over 100,000 SSA runs are still fluctuating noticeably.

Instead of performing large numbers of SSA simulations, the time evolution of higher order moments may be obtained much more efficiently from moment approximation methods. Unfortunately it is not possible to predict *a priori* how many moments are required to fully capture the output distributions of stochastic processes. And while our method can be used to evaluate arbitrary higher moments, the number of equations that need to be solved increases rapidly with the number of molecular species and the number of moments required. This comes with increased computational cost for calculating the expressions as well as for integrating the equations. Therefore it is desirable to identify whether more moments are needed to model the system, and we have provided a simple heuristic to achieve this in many cases. While this can indicate if more than two moments are needed, it will not clearly identify when sufficiently many central moments have been included in the expansion. This is a point for further future investigation. In cases where the distribution is known, but different from normal, knowledge about the distribution can be used to close the set of equations correspondingly. The estimated higher order moments may also by themselves indicate which distribution is a probable candidate for the output of a given dynamical stochastic system; we can, for example, compare obtained estimates of such moments against known values for the higher order moments for particular distributions: if e.g. the estimated skewness and kurtosis are close to zero and three, respectively, a normal distribution might be a good choice and an approximation similar to the simple linear noise approximation may be appropriate.

When the model under investigation is non-linear, lower order moments typically depend on higher order moments, and the system of equations is not closed. Closure can be achieved in several different ways. Here we have closed the system by truncating the Taylor series expansion at a selected order. Several other approaches have been previously evaluated. Chevalier et al³⁸, for example, approximated higher order moments using experimental data. Azunre et al³⁹ showed that for very small molecule numbers, using only two moments can lead to unstable results. Singh et al.¹⁶ developed a derivative-matching approach in the context of polynomial rate equations, which proved to work very well in particular for small molecule numbers. For the examples described in this paper, the proposed method shows differences in behaviour for very low molecule numbers. More specifically, for the dimerisation system, the mean trajectories become unstable below approximately 7 molecules in the system; for the Michaelis-Menten example, the system does not become unstable, but becomes less accurate for very small molecule numbers. The p53 system

trajectories for the system size as displayed in Fig. 3 start running out of phase around $t = 35$, and when the molecule number decreases, the point where the trajectories run out of phase moves forward, followed by instability in the trajectories. The LNA solution for the p53 system does not become unstable for small molecule numbers but keeps showing the expanding oscillatory behaviour. For other systems that show periodic behaviour it might be more beneficial to approximate the system in the frequency domain. We would recommend use of this method — certainly if no manual closure is attempted, but probably even then — for moderate and large numbers of molecules; a conservative estimate might be to have at least 10-20 molecules. However, for molecular abundances of less than 10 molecules we find that the numerical burden of the SSA compared to the MFK approaches is no longer prohibitive, and, again conservatively, we would consider the use of the SSA.

The present work also provides a natural framework for an inferential framework: our results above suggest that inclusion of higher order moments will lead to increased accuracy of the parameter estimates. Milner et al.⁴⁰ derived a likelihood function that included the mean and variance obtained from a moment closure method, assuming a multivariate Gaussian distribution. Kuegler⁴¹ used both the mean and the variance obtained from a moment closure method for parameter estimation by minimizing an objective function that included both the difference between the observed and estimated mean and the difference between the observed and estimated variance. This showed that more accurate parameter estimates could be obtained when the observed variance is included for parameter estimation. Other analyses have employed approximate Bayesian computation schemes⁴², and used moment-based inference employing the mean and variance to infer parameters of a bimodal system⁴³. Through recent technological advances in the field of single cell observations^{44–46}, it becomes possible to probe directly the properties of the output distributions of stochastic dynamical systems over time. The additional information about the higher order central moments that can be derived from these datasets can be exploited when higher order empirical moments are also used for parameter inference. However, while likelihoods are trivially constructed when the mean behaviour and variance estimates are available (via Gaussian assumptions), conditioning on higher-order moments typically requires some further assumptions; most easily these moments are included in maximum-entropy estimators of the probability distribution. The present approach yields these moments, however, reliably and affordably. Because of their affordability these moments also open up new ways for assessing the sensitivity of stochastic dynamical systems (as outlined in section IIID), including cases where the linear noise approximation tends to break down¹³. This includes general feedback systems where the notion of sensitivity may be particularly useful but calculation for stochastic systems is fraught with problems and numer-

ically expensive. In conclusion, the general moment expansion method described in this paper provides a flexible and valuable new tool for investigating many stochastic kinetic systems.

Appendix A: Model equations

In this section we display the model equations used for the dimerisation and Michaelis-Menten system. The complexity of the equations grows with the number of moments included, we display here the equations used for the mean, variance and co-variance when truncating after second order.

1. Dimerisation

For the dimerisation system, the equations used for the mean, variance and covariance are given by

$$\begin{aligned}\mu_{x_1} &= 2k_2x_2 - 2k_1\sigma_{x_1}^2 - 2k_1x_1(x_1 - 1) \\ \mu_{x_2} &= k_1\sigma_{x_2}^2 - k_2x_2 + k_1x_1(x_1 - 1) \\ \sigma_{x_2}^2 &= k_1y_1^2 - k_1x_1 + k_2x_2 + c_1\sigma_{x_1}^2 - 2k_2\sigma_{x_2}^2 - 2k_1\sigma_{x_1,x_2}^2 \\ &\quad + 4k_1x_1\sigma_{x_1,x_2}^2 \\ \sigma_{x_1,x_2}^2 &= 2k_1\sigma_{x_1,x_2}^2 + 2k_2\sigma_{x_2}^2 - k_1\sigma_{x_1}^2 - k_2\sigma_{x_1,x_2}^2 - 2k_1x_1^2 \\ &\quad - 2k_1x_1 + 2k_2x_2 + 2k_1\sigma_{x_1}^2 - 4k_1x_1\sigma_{x_1,x_2}^2 + 2k_1x_1\sigma_{x_1}^2 \\ \sigma_{x_1}^2 &= 4k_1\sigma_{x_1}^2 + 4k_2\sigma_{x_1,x_2}^2 + 4k_1x_1^2 - 4k_1x_1 + 4k_2x_2 + \\ &\quad 4k_1\sigma_{x_1}^2 - 8k_1x_1\sigma_{x_1}^2\end{aligned}$$

2. Michaelis-Menten System

For the Michaelis-Menten system, the equations used for the mean, variance and covariance are given by

$$\begin{aligned}\mu_{x_1} &= -k_2(x_1 + x_2 - 301) - k_1\sigma_{x_1,x_2}^2 - \\ &\quad k_1\sigma_{x_1}^2 - k_1x_1(x_1 + x_2 - 181) \\ \mu_{x_2} &= -c_3(x_1 + x_2 - 301) \\ \sigma_{x_2}^2 &= -2c_3(\sigma_{x_2}^2 + \sigma_{x_1,x_2}^2) - c_3(x_1 + x_2 - 301) \\ \sigma_{x_1,x_2}^2 &= 181k_1\sigma_{x_1,x_2}^2 - k_2\sigma_{x_2}^2 - k_2\sigma_{x_1,x_2}^2 - c_3\sigma_{x_1,x_2}^2 \\ &\quad - c_3\sigma_{x_1}^2 - k_1x_1\sigma_{x_2}^2 - 2k_1x_1\sigma_{x_1,x_2}^2 - k_1x_2\sigma_{x_1,x_2}^2 \\ \sigma_{x_1}^2 &= -(181k_1x_1 - 301k_2 + k_2x_1 + k_2x_2 \\ &\quad - k_1\sigma_{x_1,x_2}^2 - k_1\sigma_{x_1}^2 - k_1x_1^2 - 362k_1\sigma_{x_1}^2 + \\ &\quad 2k_2\sigma_{x_1,x_2}^2 + 2k_2\sigma_{x_1}^2 - k_1x_1x_2 + 2k_1x_1\sigma_{x_1,x_2}^2 \\ &\quad + 4k_1x_1\sigma_{x_1}^2 + 2k_1x_2\sigma_{x_1}^2)\end{aligned}$$

¹C. Gardiner, *Stochastic methods*, 4th ed., Springer Series in Synergetics (Springer-Verlag, Berlin, 2009).

- ²N. van Kampen, *Stochastic Processes in Physics and Chemistry* (North-Holland, 1992).
- ³D. T. Gillespie, "General method for numerically simulating stochastic time evolution of coupled chemical-reactions," *Journal of Computational Physics* **22**, 403–434 (1976).
- ⁴C. S. Gillespie, "Moment-closure approximations for mass-action models," *The Institution of Engineering and Technology* **3**, 52–58 (2009).
- ⁵C. A. Gómez-Urbe and G. C. Verghese, "Mass fluctuation kinetics: capturing stochastic effects in systems of chemical reactions through coupled mean-variance computations," *Journal of Chemical Physics* **126**, 024109–024112 (2007).
- ⁶C. Lee, K.-H. Kim, and P. Kim, "A moment closure method for stochastic reaction networks," *The Journal of Chemical Physics* **130**, 134107 (2009).
- ⁷A. Singh and J. Hespanha, "Stochastic hybrid systems for studying biochemical processes," *Phil. Trans. R. Soc. A* **368**, 4995–5011 (2010).
- ⁸A. Singh and J. Hespanha, "Lognormal moment closures for biochemical reactions," in *Decision and Control 2006 45th IEEE Conference* (2006) p. 20632068.
- ⁹J. Ruess, A. Miliadis-Argeitis, S. Summers, and J. Lygeros, "Moment estimation for chemically reacting systems by extended kalman filtering," *The Journal of Chemical Physics* **135**, 165102 (2011).
- ¹⁰J. Goutsias, "Classical versus stochastic kinetics modeling of biochemical reaction systems," *Biophysical Journal* **92**, 2350–2365 (2007).
- ¹¹B. Barzel and O. Biham, "Stochastic analysis of complex reaction networks using binomial moment equations," *Physical Review E* **86**, 031126 (2012).
- ¹²M. Komorowski, B. Finkenstaedt, C. V. Harper, and D. A. Rand, "Bayesian inference of biochemical kinetic parameters using the linear noise approximation," *BMC bioinformatics* **10** (2009).
- ¹³M. Komorowski, M. J. Costa, D. A. Rand, and M. P. H. Stumpf, "Sensitivity, robustness, and identifiability in stochastic chemical kinetics models," *PNAS* **108**, 8645–50 (2011).
- ¹⁴J. Pahle, J. Challenger, P. Mendes, and A. McKane, "Biochemical fluctuations, optimisation and the linear noise approximation," *BMC Systems Biology* **86** (2012).
- ¹⁵E. Wallace, D. Gillespie, K. Sanft, and L. Petzold, "Linear noise approximation is valid over limited times for any chemical system that is sufficiently large," *IET Syst. Biol.* **6**, 102–115 (2012).
- ¹⁶A. Singh and J. P. Hespanha, "Approximate moment dynamics for chemically reacting systems," *IEEE Transactions on Automatic Control* **56**, 414–418 (2011).
- ¹⁷V. Sotiropoulos and Y. N. Kaznessis, "Analytical derivation of moment equations in stochastic chemical kinetics," *Chemical Engineering Science* **66**, 268–277 (2011).
- ¹⁸P. Milner, C. C. Gillespie, and D. J. Wilkinson, "Moment closure approximations for stochastic kinetic models with rational rate laws," *Mathematical Biosciences* **231**, 99–104 (2011).
- ¹⁹M. Ullah and O. Wolkenhauer, "Investigating the two-moment characterisation of subcellular biochemical networks," *Journal of Theoretical Biology* **260**, 340–352 (2009).
- ²⁰R. Grima, "General method for numerically simulating stochastic time evolution of coupled chemical-reactions," *Journal of Computational Physics* **22**, 403–434 (2012).
- ²¹E. Batchelor, A. Loewer, and G. Lahav, "The ups and downs of p53: understanding protein dynamics in single cells," *Nature Reviews Cancer* **9**, 371–377 (2009).
- ²²P. Azunre, *Mass Fluctuation Kinetics: Analysis and Computation of Equilibria and Local Dynamics*, Master's thesis, Massachusetts Institute of Technology (2007).
- ²³Y. Ito and K. Uchida, "Formulas for intrinsic noise evaluation in oscillatory genetic networks," *Journal of theoretical biology* **267**, 223–234 (2010).
- ²⁴M. Komorowski, J. Miekisz, and M. P. Stumpf, "Decomposing noise in biochemical signalling systems highlights the role of protein degradation," *Biophysical Journal* **in press** (2013).
- ²⁵D. J. Wilkinson, *Stochastic Modelling for Systems Biology* (Chapman and Hall, CRC press, 2012).
- ²⁶N. Geva-Zatorsky, N. Rosenfeld, S. Itzkovitz, R. Milo, A. Sigal, E. Dekel, T. Yarnitzky, Y. Liron, P. Polak, G. Lahav, and U. Alon, "Oscillations and variability in the p53 system," *Molecular Systems Biology* **2**, 006.0033 (2006).
- ²⁷P. D. Kirk, T. Toni, and M. P. Stumpf, "Parameter inference for biochemical systems that undergo a hopf bifurcation," *Biophysical Journal* **95**, 540549 (2008).
- ²⁸A. Saltelli, S. Tarantola, F. Campolongo, and M. Ratto, *Sensitivity analysis in practice* (John Wiley & Sons Ltd., 2004).
- ²⁹K. K. Erguler and M. P. H. Stumpf, "Practical limits for reverse engineering of dynamical systems: a statistical analysis of sensitivity and parameter inferability in systems biology models," *Molecular bioSystems* **7**, 1593–1602 (2011).
- ³⁰R. Gunawan, Y. Cao, L. Petzold, and F. J. Doyle, "Sensitivity analysis of discrete stochastic systems," *Biophysical Journal* **88**, 2530–2540 (2005).
- ³¹S. Plyasunov and A. P. Arkin, "Efficient stochastic sensitivity analysis of discrete event systems," *Journal of Computational Physics* **221**, 724–738 (2007).
- ³²M. Komorowski, J. Zurauskiene, and M. P. H. Stumpf, "StochSens—matlab package for sensitivity analysis of stochastic chemical systems," *Bioinformatics (Oxford, England)* **28**, 731–733 (2012).
- ³³P. W. Sheppard, M. Rathinam, and M. Khammash, "SPSens: A software package for stochastic parameter sensitivity analysis of biochemical reaction networks," *Bioinformatics (Oxford, England)* , – (2012).
- ³⁴A. Varma, M. Morbidelli, and H. H. Wu, *Parametric sensitivity in chemical systems* (Cambridge, U.K. ; New York, NY : Cambridge University Press, 1999).
- ³⁵R. Serban and A. C. Hindmarsh, "CVODES: An ODE solver with sensitivity analysis capabilities," in *Preprint UCRL-JP-200039, Lawrence Livermore National Laboratory* (2003).
- ³⁶A. Gelman, C. J.B., H. Stern, and D. Rubin, *Bayesian Data Analysis*, 2nd ed. (Chapman & Hall/CRC, 2003).
- ³⁷T. Toni, D. Welch, N. Strelkowa, D. Ipsen, and M. Stumpf, "Approximate bayesian computation scheme for parameter inference and model selection in dynamical systems," *J.Roy.Soc. Interface* **6**, 187–202 (2009).
- ³⁸M. W. Chevalier and H. El-Samad, "A data-integrated method for analyzing stochastic biochemical networks," *The Journal of Chemical Physics* , 214110 (2011).
- ³⁹P. Azunre, C. Gómez-Urbe, and G. Verghese, "Mass fluctuation kinetics: analysis and computation of equilibria and local dynamics," *IET Systems Biology* **5**, 325–335 (2011).
- ⁴⁰P. Milner, C. C. Gillespie, and D. J. Wilkinson, "Moment closure based parameter inference of stochastic kinetic models," *Statistics and Computing* **231**, 99–104 (2012).
- ⁴¹P. Kuegler, "Moment fitting for parameter inference in repeatedly and partially observed stochastic biological models," *Plos one* **7**, 1–15 (2012).
- ⁴²T. Toni, Y.-i. Ozaki, P. Kirk, S. Kuroda, and M. P. H. Stumpf, "Elucidating the in vivo phosphorylation dynamics of the ERK MAP kinase using quantitative proteomics data and Bayesian model selection," *Molecular BioSystems* **8**, 1921–1929 (2012).
- ⁴³C. Zechner, J. Ruess, P. Krenn, S. Pelet, M. Peter, J. Lygeros, and H. Koepl, "Moment-based inference predicts bimodality in transient gene expression," *PNAS* (2012).
- ⁴⁴A. Salehi-Reyhani, J. Kaplinsky, E. Burgin, M. Novakova, R. H. T. Andrew J. deMello, P. Parker, M. A. A. Neil, O. Ces, P. French, K. R. Willison, and D. Klug, "A first step towards practical single cell proteomics: a microfluidic antibody capture chip with tirf detection," *Lab on a Chip* , 1256–1261 (2011).
- ⁴⁵Y. Lin, R. Trouillon, G. Safina, and A. G. Ewing, "Chemical analysis of single cells," *Analytical Chemistry* **83**, 43694392 (2011).
- ⁴⁶Y.-i. Ozaki, S. Uda, T. H. Saito, J. Chung, H. Kubota, and S. Kuroda, "A quantitative image cytometry technique for time

series or population analyses of signaling networks,” PLoS One **5**, e9955 (2010).